

1 **EXACTLY SOLVING SPARSE RATIONAL LINEAR SYSTEMS VIA**
2 **ROUNDOFF-ERROR-FREE CHOLESKY FACTORIZATIONS***

3 CHRISTOPHER LOURENCO [†] AND ERICK MORENO-CENTENO [‡]

4 **Abstract.**

5 Exactly solving sparse symmetric positive definite (SPD) linear systems is a key problem in math-
6 ematics, engineering, and computer science. This paper derives two new sparse roundoff-error-free
7 (REF) Cholesky factorization algorithms which exactly solve sparse SPD linear systems $A\mathbf{x} = \mathbf{b}$,
8 where $A \in \mathbb{Q}^{n \times n}$ and $\mathbf{x}, \mathbf{b} \in \mathbb{Q}^{n \times p}$. The key properties of these factorizations are: (1) they ex-
9 clusively use integer-arithmetic and (2) in the bit-complexity model, they solve the linear system
10 $A\mathbf{x} = \mathbf{b}$ in time proportional to the cost of the integer-arithmetic operations. Namely, the overhead
11 related to data-structures and ancillary operations (those not strictly required to perform the fac-
12 torization) is subsumed by the cost of the integer-arithmetic operations that are essential/intrinsic
13 to the factorization. Notably, to-date our algorithms are the only exact algorithm for solving SPD
14 linear systems with this asymptotically efficient complexity bound. Computationally, we show that
15 the novel factorizations are faster than both sparse rational-arithmetic LDL and sparse exact LU
16 factorization. Altogether, the derived sparse REF Cholesky factorizations present a framework to
17 solve any rational SPD linear system exactly and efficiently.

18 **Key words.** Cholesky Factorization, Exact Factorization, Sparse Matrix Algorithms

19 **AMS subject classifications.** 15A23, 65F50, 65F05, 65G50

20 **1. Introduction.** Solving sparse symmetric positive definite (SPD) systems of
21 linear equations (SLEs) is a fundamental problem in mathematics, computer science,
22 and operations research. In general, SLEs with an SPD input matrix are solved via
23 Cholesky factorization, where the input matrix, A , is factored into the product of a
24 lower triangular matrix, L , and its transpose; that is $A = LL^T$. For SPD matrices,
25 (floating-point) Cholesky factorization is normwise backward stable [45, 35], meaning
26 that the relative residual norm $\|A\mathbf{x} - \mathbf{b}\|/(\|A\|\|\mathbf{x}\|)$ is guaranteed to be close to
27 machine precision. However, despite strong residual error bounds, as Golub & Van
28 Loan point out, “small residuals do not imply high [solution] accuracy” [21]. Indeed,
29 Cholesky factorization may fail if the input is highly ill-conditioned [45, 35, 27, 44]
30 and Cholesky is guaranteed to fail if the singular values of the input matrix are too
31 small [10]. That said, it is unclear how many real world matrices may cause Cholesky
32 factorization to fail, and the correctness of commercial solvers is typically taken for
33 granted; thus, there is currently no exact algorithm specifically tuned to provide exact
34 solutions to SPD linear systems. Developing such algorithms is the core contribution
35 of this paper.

36 **1.1. Contributions.** To address the issues outlined above, this paper derives
37 two new sparse roundoff-error-free (REF) Cholesky factorization algorithms. The
38 two new factorizations extend the (dense) REF Cholesky factorization [16] and the
39 sparse left-looking integer-preserving (SLIP) LU factorization [33]. Specifically, we
40 extend the REF Cholesky factorization to the sparse case by deriving both an up-

*Submitted to the editors July 15, 2021.

Funding: This work was supported in part by the National Science Foundation under Grant No, OAC-1835499. In addition, the first author was also partially supported by the Texas A&M University Graduate Merit Fellowship and the US Naval Academy Junior NARC.

[†]Department of Mathematics, United States Naval Academy, Annapolis, Maryland 21402 (lourenco@usna.edu)

[‡]Department of Industrial and Systems Engineering, Texas A&M University, College Station, Texas 77843 (emc@tamu.edu)

41 looking and left-looking variant (two of the most common types used in practice
 42 today). In addition, we modify the key subroutine within the SLIP LU factorization
 43 to exploit symmetry, skip ancillary operations, and solve both lower triangular and
 44 lower trapezoidal linear systems in order to derive two new algorithms to compute
 45 the aforesaid factorizations in a sparse manner. These two new sparse REF Cholesky
 46 factorizations are unique in that, unlike traditional Cholesky factorization, computing
 47 the REF factor L does not require the usage of square root functions; thus they may
 48 be applied to *any* rational SPD input matrix $A \in \mathbb{Q}^{n \times n}$.

49 In addition, we analyze the computational complexity of both factorizations.
 50 Specifically, we show that, using a bit-complexity model, both of our algorithms solve
 51 the SPD sparse linear system $A\mathbf{x} = \mathbf{b}$ in time proportional to the cost of the integer-
 52 arithmetic operations. This means that all of the overhead related to data structures
 53 and ancillary operations is asymptotically dominated by the cost of the indispens-
 54 able arithmetic operations to compute the factorization. Notably, this property is
 55 not trivial in the sparse case, as currently, to the best of our knowledge, there are
 56 only three sparse factorizations to achieve this asymptotically efficient complexity
 57 bound: Gilbert and Peierls' (floating-point) left-looking LU factorization [19], Liu's
 58 (floating-point) up-looking Cholesky factorization [31], and (exact) left-looking SLIP
 59 LU factorization [33]. It is important to note that many classic sparse matrix algo-
 60 rithms can not meet such asymptotic efficiency as their run times may be dominated
 61 by ancillary operations such as data manipulation, access, and/or allocations; notably
 62 this includes algorithms as fundamental (and seemingly trivial) as sparse matrix mul-
 63 tiplication [8].

64 Next, we perform a computational study which shows that the derived factoriza-
 65 tions are dramatically faster than exact unsymmetric LU factorization and rational-
 66 arithmetic LDL factorization. Finally, we benchmark the relative forward error of
 67 MATLAB sparse backslash, showing that state-of-the-art floating-point Cholesky
 68 factorization may fail for a modest subset of real-world instances; thereby demon-
 69 strating the need for our algorithms to exactly solve SPD linear systems. The
 70 code associated with the sparse REF Cholesky factorizations is freely available at
 71 https://github.com/clouren/IP_Chol.

72 **1.2. Organization.** This paper is organized as follows. Section 2 gives an
 73 overview of the background required for this paper. Section 3 derives the up-looking
 74 and left-looking REF Cholesky factorizations. Section 4 analyzes the computational
 75 complexity of both factorizations. Section 5 computationally compares the derived
 76 Cholesky factorizations with exact unsymmetric factorization and rational LDL fac-
 77 torization on real-world SPD instances for run time and MATLAB sparse backslash
 78 for accuracy. Finally, Section 6 concludes the paper.

79 **2. Background.** This section describes the theoretical foundation of the new
 80 factorizations and is organized as follows. Subsection 2.1 discusses Integer-preserving
 81 Gaussian elimination, the basis of our algorithms. Subsections 2.2, 2.3, and 2.4 dis-
 82 cuss (floating point) up-looking Cholesky, left-looking Cholesky, and LDL factoriza-
 83 tion, respectively. Finally, subsections 2.5 and 2.6 review the (dense) REF Cholesky
 84 factorization and the SLIP LU factorization respectively. Note that, for the sake of
 85 simplicity, the descriptions in this section assume no row or column permutations are
 86 applied to the matrix A .

87 **2.1. Integer-Preserving Gaussian Elimination.** Integer-preserving Gaus-
 88 sian elimination (IPGE) is an elimination process for solving a system of linear equa-

89 tions $A\mathbf{x} = \mathbf{b}$. Given a full rank matrix $A \in \mathbb{Z}^{n \times n}$ and right hand side vector \mathbf{b} ,
 90 denote the k th iteration IPGE matrix as $A^{(k)}$ for $k = 0, \dots, n$ (where $A^{(0)} = A$).
 91 Let $a_{i,j}^k$ and ρ^k denote the individual entries of $A^{(k)}$ and the k th pivot element for
 92 $1 \leq i \leq n$, $1 \leq j \leq n$, and $0 \leq k \leq n$ (with $\rho^0 = 1$), respectively. Then, at iteration
 93 k , the IPGE algorithm computes the entries $a_{i,j}^k$ as follows:

$$94 \quad (1) \quad a_{i,j}^k = \begin{cases} a_{i,j}^{k-1} & \text{if } i = k, \\ \frac{\rho^k a_{i,j}^{k-1} - a_{k,j}^{k-1} a_{i,k}^{k-1}}{\rho^{k-1}} & \text{otherwise} \end{cases}$$

95 Note that Equation (1) differs from traditional Gaussian elimination only in the
 96 denominator (in traditional Gaussian elimination, the division is by the current pivot,
 97 ρ^k , instead of the previous pivot). This seemingly minor modification leads to two
 98 key properties summarized below.

99 LEMMA 2.1. *The divisions in Equation (1) are integral [14, 4, 36].*

100 Lemma 2.1 has been proven in three different manners.

101 Both Edmonds and Bareiss prove the integrality of the IPGE entries by showing
 102 (each in a different way) that each IPGE entry is a subdeterminant of A . Specifically,
 103 by using row expansions, Edmonds [14] shows that each IPGE entry, $a_{i,j}^k$, is equiva-
 104 lent to a corresponding subdeterminant of A . Bareiss [4] proves the integrality of the
 105 IPGE entries using Sylvester's identity. Specifically, he shows that a block triangu-
 106 lar partitioning of A can be manipulated to take the form of Sylvester's identity
 107 [37], and then, uses various sub determinant properties to derive the IPGE equation.
 108 Montante-Pardo and Mendez-Cavazos [36] aim to prove the integrality of IPGE by
 109 induction; while this proof is good for building intuition, it falls short as they do not
 110 show the induction hypothesis to be true for a general k .

111 LEMMA 2.2. *Let $\sigma = \max_{i,j} |a_{i,j}^0|$. Given an SPD matrix A , the maximum bit length
 112 to store any IPGE entry, denoted β_{max} , is upper-bounded polynomially as follows [33]:*

$$113 \quad \beta_{max} \leq \lceil n \log(\sigma) \rceil.$$

114 Lemma 2.2 is proved as follows. Since each IPGE entry is a subdeterminant of
 115 A , its worst case bit-length can be upper-bounded polynomially. This fact was first
 116 shown in [5] and further revised in [29, 16]. The current tightest bound known for
 117 SPD matrices is the one presented in Lemma 2.2 which was derived by [33] by using
 118 Hadamard's bound for SPD matrices [23, 25]. Note that though this bound is currently
 119 the best known, it is pessimistic, as it is tight if and only if the matrix A is diago-
 120 nal. That said, Lemma 2.2 is the key property making IPGE preferable to rational-
 121 arithmetic Gaussian elimination. Specifically, IPGE achieves gratis the polynomial
 122 bound on the entries; in contrast, the bit-lengths of the numerators/denominators in
 123 rational-arithmetic algorithms grow exponentially without the usage of ancillary and
 124 computationally expensive greatest common divisor operations [43, 42].

125 **2.2. Traditional Up-Looking Cholesky Factorization.** Given an SPD ma-
 126 trix A , the up-looking Cholesky factorization computes the sparse Cholesky factor-
 127 ization, $A = LL^T$, one row at a time. Graphically, Figure 1 illustrates the data access
 128 pattern of up-looking Cholesky factorization. Specifically, at iteration k , up-looking
 129 Cholesky accesses the $(k-1) \times (k-1)$ completed portion of L (denoted $L(1:k-1, 1:k-1)$)
 130 and the first k entries of the k th row of A (denoted $A(k, 1:k)$) and computes the k th
 131 row of L (thereby the k th column of L^T). Up-looking Cholesky factorization computes

132 the k th row of L by solving a lower triangular linear system and then computing a
 133 square root. Formally, the solutions of the equations below yield the k th row of L ,
 134 where \mathbf{x} is the first $k - 1$ entries of this row and $l_{k,k}$ is the k th entry of the row.

$$135 \quad L(1 : k - 1, 1 : k - 1)\mathbf{x} = A(k : 1 : k - 1)^T,$$

$$136 \quad l_{k,k} = \sqrt{a_{k,k}^0 - \mathbf{x}^T \mathbf{x}}.$$

138 Assuming that A has a fixed row/column ordering, the up-looking Cholesky fac-
 139 torization computes L in time proportional to only the number of intrinsic floating-
 140 point operations; meaning that the cost of these indispensable floating-point opera-
 141 tions asymptotically dominates the cost of all data access, manipulation, and ancil-
 142 lary operations. As a result, this algorithm's computational complexity is directly
 143 proportional to its arithmetic work. Up-looking Cholesky factorization was derived
 144 in [40, 30, 3]; however, the modern version, which achieved the aforesaid asymptotic
 145 efficiency, was developed by [31].

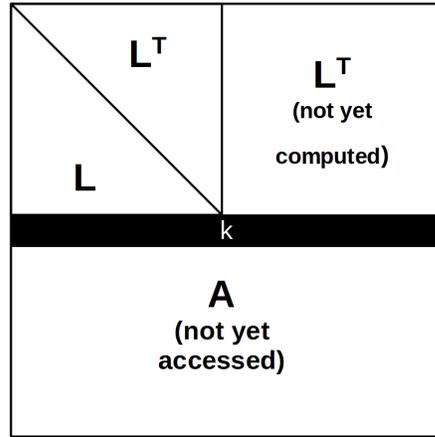


Fig. 1: Up-looking Cholesky Factorization

146 **2.3. Traditional Left-Looking Cholesky Factorization.** Given an SPD ma-
 147 trix A , left-looking Cholesky factorization computes the sparse Cholesky factorization,
 148 $A = LL^T$, one column at a time. Graphically, Figure 2 illustrates the data access
 149 pattern of left-looking Cholesky factorization. Specifically, at iteration k , left-looking
 150 Cholesky accesses rows k, \dots, n of the first $k - 1$ computed columns of L (denoted
 151 $L(k : n, 1 : k - 1)$), the k th row of L (denoted $L(k, 1 : k - 1)$), the k th column of L
 152 (denoted $L(1 : k - 1, k)$) and the last $n - k$ entries of the k th column of A (denoted
 153 $A(k + 1 : n, k)$). Formally, the equations below give the k th column of the matrix L .

$$154 \quad l_{k,k} = \sqrt{a_{k,k} - L(k, 1 : k - 1) \cdot L(k, 1 : k - 1)^T},$$

$$155 \quad L(k + 1 : n, k) = (A(k + 1 : n, k) - L(k : n, 1 : k - 1)L(1 : k - 1, k))/l_{k,k}.$$

157 Notice that, unlike up-looking Cholesky, left-looking Cholesky factorization is
 158 based on performing a matrix vector product (specifically, computing $L(k : n, 1 :$

159 $k - 1) \cdot L(1 : k - 1, k)$). To the best of our knowledge, it is not claimed that this
 160 factorization meets the same asymptotically efficient complexity bound of up-looking
 161 Cholesky. Historically, left-looking Cholesky preceded the up-looking variants and can
 162 be attributed to [38, 40, 15, 18]. We conclude this subsection by noting that both left-
 163 looking and up-looking Cholesky factorization is used in the state-of-the-art package
 164 CHOLMOD [7] for solving sparse SPD systems. Specifically, up-looking Cholesky is
 165 utilized for very sparse input matrices while left-looking is used for all other matrices.

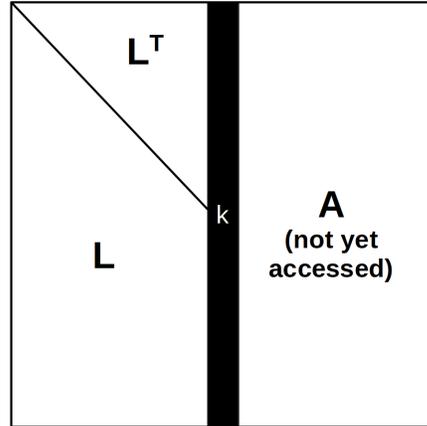


Fig. 2: Left-looking Cholesky Factorization

166 **2.4. LDL Factorization.** Given an SPD matrix A , LDL factorization computes
 167 the sparse symmetric factorization, $A = LDL^T$, where L is a unit lower triangular
 168 matrix and D is a diagonal matrix. In general, there are two variants of LDL fac-
 169 torization, the Bunch-Parlett LDL factorization (described below) [6] and the LDL
 170 factorization with a block diagonal \mathbf{D} matrix. This second method, detailed in [12],
 171 is outside the scope of this paper as it is generally applied to symmetric indefinite
 172 linear systems.

173 The Bunch-Parlett LDL factorization exploits the symmetry of an SPD input
 174 matrix; however, in contrast to Cholesky factorization, avoids the use of square roots.
 175 Thus, LDL factorization can be used with rational-arithmetic to exactly solve SPD
 176 linear systems. Specifically, the LDL factorization generalizes Cholesky factorization
 177 since the Cholesky factor L' can be obtained from the LDL factors L and D as
 178 $L' = L\sqrt{D}$ [21].

179 **2.5. Dense REF Cholesky Factorization.** The REF Cholesky factorization
 180 solves a dense SPD system of linear equations exclusively in integer arithmetic [16].
 181 Formally, the SPD matrix $A \in \mathbb{Z}^{n \times n}$ is factored as $A = (L\sqrt{D})(L\sqrt{D})^T = LDL^T$
 182 where $D = \text{diag}(\rho^0 \rho^1, \rho^1 \rho^2, \dots, \rho^{n-1} \rho^n)^{-1}$ and the entries in L are IPGE entries
 183 (specifically, $l_{i,j} = a_{i,j}^{\min(i,j)-1}$). Note that this factorization is indeed an integral
 184 Cholesky factorization, as the matrix D is never explicitly constructed or used when
 185 solving SPD linear systems.

186 **2.6. SLIP LU.** The SLIP LU factorization is a left-looking sparse factorization
 187 that computes the exact solution of a sparse linear system $A\mathbf{x} = \mathbf{b}$ solely using integer

188 arithmetic. Based on the roundoff-error-free (REF) LU factorization [16], SLIP LU
 189 computes the factorization $A = LDU$ one column at a time. At iteration k , SLIP LU
 190 accesses the first $k - 1$ columns of L , denoted $L_L^{(k-1)}$, and the first $k - 1$ columns of
 191 D augmented by the last $n - k + 1$ columns of $(1/\rho^{k-1})I$, denoted $D_L^{(k-1)}$, and solves
 192 the following linear system:

$$193 \quad L_L^{(k-1)} D_L^{(k-1)} \mathbf{x} = A(:, k).$$

194 The solution \mathbf{x} to this linear system yields the (integral) k th column of the L and
 195 U matrices. However, this system can not be solved directly with a traditional lower
 196 triangular solve algorithm. Instead, it is solved in two phases: symbolic and numeric.

197 Symbolically, the nonzero pattern of the k th column of L and U is computed via
 198 a graph algorithm. Namely, this algorithm performs a sequence of depth first searches
 199 (DFS) on a graph G_L of n nodes, where a directed edge (i, j) exists if $l_{j,i} \neq 0$. The
 200 nonzero pattern, denoted \mathcal{X} , is given as the indices of the set of nodes reachable, via
 201 DFS, from each node j in which $a_{j,k}^0 \neq 0$. This algorithm is referred to as obtaining the
 202 reach of the k th column of A on the graph of L and is denoted $\mathcal{X} = Reach_{G_L}(A(:, k))$.
 203 After \mathcal{X} is computed, it is sorted prior to numeric computation.

204 Numerically, this system is solved via a sequence of two operations. The first,
 205 referred to as a History Update, involves multiplying the entry x_j by the current pivot
 206 and dividing it by a previous pivot which is an essential operation in order to exploit
 207 sparsity in IPGE. The second operation is referred to as an IPGE Update (Equation
 208 (1)). For completeness, this algorithm is given below as Algorithm 1.

Algorithm 1 Sparse REF Triangular Solve

```

1:  $x = A(:, k)$ 
2:  $\mathcal{X} = Reach_{G_L}(A(:, k))$ 
3: sort( $\mathcal{X}$ )
4: Initialize history vector:  $h_j = 0 \forall j \in \mathcal{X}$ 
5: for  $j \in \mathcal{X}$  do
6:   if  $j < k$  then
7:     if  $h_j < j - 1$  then
8:       History update:  $x_j = \frac{x_j \rho^{j-1}}{\rho^{h_j}}$ 
9:     end if
10:    for  $i > j$  and  $l_{i,j} \neq 0$  do
11:      if  $h_i < j - 1$  then
12:        History update:  $x_i = \frac{x_i \rho^{j-1}}{\rho^{h_i}}$ 
13:      end if
14:      IPGE update:  $x_i = \frac{\rho^j x_i - l_{i,j} x_j}{\rho^{j-1}}$ 
15:      History vector update:  $h_i = j$ 
16:    end for
17:  else
18:    if  $h_j < k - 1$  then
19:      History update:  $x_j = \frac{x_j \rho^{k-1}}{\rho^{h_j}}$ 
20:    end if
21:  end if
22: end for

```

209 Applied n times, Algorithm 1 factors the matrix $A \in \mathbb{Z}^{n \times n}$ by performing only

210 the *indispensable* number of full-precision operations. This is in contrast to rational
 211 factorization, for example, which must perform ancillary greatest common divisor
 212 operations to limit bit-growth. To date, SLIP LU is the only sparse exact factorization
 213 method for solving unsymmetric linear systems in time proportional to the arithmetic
 214 work using the bit-complexity model [33] subject to the same caveat as (floating-point)
 215 up-looking Cholesky factorization, namely that the row/column ordering is fixed.

216 **3. Sparse REF Cholesky Factorizations.** This section formally defines the
 217 two new sparse REF Cholesky factorizations, which can be applied to any rational
 218 SPD input matrix. Both factorizations are computed by repeatedly solving sparse
 219 lower triangular linear systems. Thus, in order to derive each factorization, we derive
 220 two new REF triangular solve algorithms which modify Algorithm 1 to exploit sym-
 221 metry. Subsection 3.1 derives Algorithm 2, the basis of the up-looking REF Cholesky
 222 factorization, and Subsection 3.2 derives Algorithm 3, the basis of the left-looking
 223 REF Cholesky factorization. Notice that, unlike the floating-point case, the left-
 224 looking REF Cholesky factorization requires solving lower triangular systems in order
 225 to guarantee integrality when computing the factorization (due to the interrelation of
 226 consecutive pivots in IPGE). Henceforth, all matrices are assumed to be SPD, and
 227 thus, $\rho^k = a_{k,k}^{k-1}$ for all $1 \leq k \leq n$ with $\rho^0 = 1$.

228 **3.1. Up-Looking REF Cholesky Factorization.** Given an SPD matrix A ,
 229 the up-looking REF Cholesky factorization computes a sparse version of the REF
 230 Cholesky factorization, $A = LDL^T$, one row at a time. At iteration k , the up-looking
 231 REF Cholesky factorization accesses the first $k - 1$ rows and columns of L and D
 232 and computes the k th row of the REF Cholesky L matrix. Prior to presenting this
 233 factorization, we define the following notation:

234 **DEFINITION 3.1.** Let $L_U^{(k)}$ and $D_U^{(k)}$ denote the k th up-looking REF L and D ma-
 235 trices. Specifically, $L_U^{(k)}$ and $D_U^{(k)}$ are the first k completed rows and columns of L
 236 and D , respectively.

237 Theorem 3.2 gives the up-looking REF Cholesky factorization.

238 **THEOREM 3.2.** The REF Cholesky factorization can be obtained in an up-looking
 239 fashion as follows. Initialize $\rho^1 = a_{1,1}^0$. Then, for $k = 2, \dots, n$, solve Equation (2)
 240 and subsequently apply Equation (3).

$$241 \quad (2) \quad L_U^{(k-1)} D_U^{(k-1)} \mathbf{x} = A(1 : k - 1, k),$$

242

$$243 \quad (3) \quad \rho^k = \rho^{k-1} \left(a_{k,k}^0 - \sum_{i=1}^{k-1} \frac{l_{k,i}^2}{\rho^i \rho^{i-1}} \right)$$

244 Prior to proving this theorem, we present and prove the following Lemma.

245 **LEMMA 3.3.** For any $k \in \{2, \dots, n\}$ suppose that $L_U^{(k-1)} D_U^{(k-1)}$ is correct. Then
 246 solving Equations (2) and (3) yields the k th row of L . Specifically, $\mathbf{x} = L(k, 1 : k - 1)^T$
 247 and $\rho^k = L(k, k)$.

248 *Proof.* To prove this Lemma, we will first derive an expression for entries $1 : k - 1$
 249 of the k th row of L and the pivot element using the first k rows and columns of the
 250 REF Cholesky factorization, that is L_U^k and D_U^k (notice that the first $k - 1$ rows
 251 and columns of this factor are $L_U^{k-1} D_U^{k-1}$). We will then show that these derived

expressions are equivalent to Equations (2) and (3), respectively. The REF Cholesky factorization of $A(1 : k, 1 : k)$ is given as:

$$L_U^{(k)} D_U^{(k)} (L_U^{(k)})^T = A(1 : k, 1 : k)$$

Let \mathbf{l}_{k-1} and \mathbf{a}_{k-1} denote entries $1 : k - 1$ of the k th column of $(L_U^{(k)})^T$ and A , respectively. Then, a 2×2 block decomposition of the above factorization is:

$$\begin{bmatrix} L_U^{(k-1)} D_U^{(k-1)} & 0 \\ \mathbf{l}_{k-1}^T D_U^{(k-1)} & L(k, k) D(k, k) \end{bmatrix} \begin{bmatrix} (L_U^{(k-1)})^T & \mathbf{l}_{k-1} \\ 0 & L(k, k) \end{bmatrix} = \begin{bmatrix} A(1 : k - 1, 1 : k - 1) & \mathbf{a}_{k-1} \\ \mathbf{a}_{k-1}^T & a_{k,k}^0 \end{bmatrix}$$

Performing the matrix multiplication in the above decomposition, yields the following two equations:

$$(4) \quad L_U^{(k-1)} D_U^{(k-1)} \mathbf{l}_{k-1} = \mathbf{a}_{k-1}$$

261

$$(5) \quad \mathbf{l}_{k-1}^T D_U^{(k-1)} \mathbf{l}_{k-1} + L(k, k) D(k, k) L(k, k) = a_{k,k}^0$$

Notice that Equation (4) is identical to Equation (2) thus $\mathbf{x} = \mathbf{l}_{k-1}$.

Likewise, recall that $D = \text{diag}(\rho^0 \rho^1, \rho^1 \rho^2, \dots, \rho^{n-1} \rho^n)^{-1}$; thus $D(k, k) = \frac{1}{\rho^k \rho^{k-1}}$,

$L(k, k) = \rho^k$ and

$$\mathbf{l}_{k-1}^T D_U^{(k-1)} \mathbf{l}_{k-1} = \sum_{i=1}^{k-1} \frac{L(k, i)^2}{\rho^i \rho^{i-1}}$$

Substituting the above expressions into Equation (5) and simplifying we obtain:

$$\sum_{i=1}^{k-1} \frac{l_{k,i}^2}{\rho^i \rho^{i-1}} + \frac{\rho^k}{\rho^{k-1}} = a_{k,k}^0$$

Solving for ρ^k we obtain:

$$\rho^k = \rho^{k-1} \left(a_{k,k}^0 - \sum_{i=1}^{k-1} \frac{l_{k,i}^2}{\rho^i \rho^{i-1}} \right)$$

The above equation is the same as Equation (3), which completes the proof. \square

Now we prove the correctness of Theorem 3.2.

Proof. This proof will use induction on k to show that solving Equations (2) and (3), for $k = 1, \dots, n$, correctly obtains the REF Cholesky factorization of A in an up-looking fashion.

Base Case: $k = 1$. At step 1, we initialize $\rho^1 = l_{1,1} = a_{1,1}^0$. This is the correct first row of L .

Induction Hypothesis: For all $j < k$ assume that the solution of Equations (2) and (3) give the correct k th row of L . Then, for row k , the induction hypothesis tells us that $L^{(k-1)} D^{(k-1)}$ is correct. Lemma 3.3 showed that since $L^{(k-1)} D^{(k-1)}$ is correct, then Equations (2) and (3) correctly compute the k th row of L . Since this is true for an arbitrary k it is true for all k . \square

283 Theorem 3.2 shows how to compute the up-looking REF Cholesky factorization
 284 by solving a sequence of lower triangular systems (Equation 2). However, note that,
 285 solving this system directly via a traditional (floating-point) lower triangular solve
 286 algorithm would introduce roundoff errors (this is easily noticed in Equation (3)).
 287 In order to solve these systems without roundoff errors, one could simply use SLIP
 288 LU's triangular solve algorithm (i.e., Algorithm 1). However, while avoiding roundoff
 289 errors, this approach would be inefficient for the following reasons: (1) Algorithm
 290 1 assumes that the system is lower trapezoidal not strictly lower triangular (since
 291 it is for left-looking LU), (2) Algorithm 1 performs its symbolic analysis without
 292 exploiting the special structure of SPD instances, and (3) Algorithm 1 computes the
 293 entire row/column which would double the work for Cholesky factorization. Thus,
 294 we modify Algorithm 1 to both exploit the properties of SPD matrices and symmetry
 295 (while maintaining its integer-preserving property) as follows:

296 First, we modify the symbolic analysis. Since Equation (2) is a lower triangular
 297 system of the form $L^{(k-1)}D^{(k-1)}\mathbf{x} = A(1 : k-1, k)$, the conditions for a nonzero
 298 entry in \mathbf{x} are identical to those in the unsymmetric case, specifically: x_j is nonzero
 299 if $a_{j,k}$ is nonzero or if both $l_{i,k}$ and $l_{j,i}$ are nonzero for some i . Thus, it is valid to
 300 compute the nonzero pattern of \mathbf{x} as the reach of the k th column of A on the graph of
 301 L . However, since the input matrix is SPD, the graph reachability can be computed
 302 more efficiently than in the unsymmetric case. Specifically, some of the conditions for
 303 a nonzero x_j discussed above are redundant for an SPD matrix [39]; thus, instead of
 304 analyzing the graph of L as described in Subsection 2.6, a simpler graph, referred to
 305 as the elimination tree, can be used. As its name suggests, the elimination tree of
 306 the $(k-1) \times (k-1)$ matrix $L_U^{(k-1)}$ is an in-tree of $k-1$ nodes (i.e., a directed tree with
 307 edges directed towards the root node), where node j is the parent of node i if the first
 308 off-diagonal nonzero in column i appears in row j . Then the nonzero pattern of \mathbf{x} ,
 309 denoted \mathcal{X} , is given as the set of indices of all nodes reachable while traversing (up)
 310 the elimination tree beginning at each node j in which $a_{j,k} \neq 0$ [30]. Consequently,
 311 this algorithm is referred to as obtaining the reach of the k th column of A on the
 312 elimination tree of A , and is denoted as $\mathcal{X} = \text{Reach}_{T_A}(A(1 : k-1, k))$. Note that
 313 this algorithm is identical to the symbolic analysis step performed in traditional up-
 314 looking Cholesky factorization [31] due to the fact that the nonzero patterns of the
 315 REF and traditional Cholesky factorizations are identical [17].

316 Second, we modify the numeric routines to improve efficiency as follows. The
 317 matrices $L_U^{(k-1)}$ and $D_U^{(k-1)}$ are of dimension $(k-1) \times (k-1)$, thus this system is indeed
 318 lower triangular and not lower trapezoidal (as is the system solved by Algorithm (1)).
 319 As a result, the up-looking triangular solve does not have to determine whether a
 320 nonzero lies above or below the current pivot element (as is done in line 6 of Algorithm
 321 1). In addition, notice that computing Equation (3) directly is inefficient, as it requires
 322 access to previously computed entries in the k th row. Moreover, computing the pivot
 323 directly as shown is not guaranteed to preserve integrality. To remedy this issue, we
 324 note that Equation (3) is of the same format as Equation (8) of [33]; thus, as argued
 325 in that paper, it can be computed via a sequence of IPGE and History updates. To
 326 perform these operations, we update the value of ρ^k after obtaining the final value of
 327 each x_j (since each $x_j = l_{k,j}$). These numeric modifications result in performing half
 328 the number of operations as the unsymmetric variant. Taken together, the symbolic
 329 and numeric improvements lead to Algorithm 2, an efficient and exact method to
 330 compute the k th row of the integer preserving Cholesky factor L .

Algorithm 2 Symmetric Sparse REF Triangular Solve (Up-Looking)

```

1: Initialization:  $x = A(1 : k - 1, k)$ 
2: Nonzero Pattern:  $\mathcal{X} = Reach_{T_L}(A(1 : k - 1, k))$ 
3: sort( $\mathcal{X}$ )
4: Initialize history vector:  $h_j = 0 \forall j \in \mathcal{X}$ 
5: Initialize pivot history value:  $h' = 0$ 
6: for  $j \in \mathcal{X}: j \neq k$  do
7:   if  $h_j < j - 1$  then
8:     History Update:  $x_j = \frac{x_j \rho^{j-1}}{\rho^{h_j}}$ 
9:   end if
10:  for  $i > j$  do
11:    if  $h_i < j - 1$  then
12:      History Update:  $x_i = \frac{x_i \rho^{j-1}}{\rho^{h_i}}$ 
13:    end if
14:    IPGE Update:  $x_i = \frac{\rho^j x_i - l_{i,j} x_j}{\rho^{j-1}}$ 
15:    History Vector Update:  $h_i = j$ 
16:  end for
17:  if  $h' < j - 1$  then
18:    History Update:  $\rho^k = \frac{\rho^k \rho^{j-1}}{\rho^{h'}}$ 
19:  end if
20:  IPGE Update:  $\rho^k = \frac{\rho^j \rho^k - x_j x_j}{\rho^{j-1}}$ 
21:  History Vector Update:  $h' = j$ 
22: end for
23: if  $h' < k - 1$  then
24:  History Update:  $\rho^k = \frac{\rho^k \rho^{k-1}}{\rho^{h_k}}$ 
25: end if

```

331 **3.2. Left-Looking REF Cholesky Factorization.** Given an SPD matrix A ,
332 the left-looking REF Cholesky factorization computes a sparse version of the REF
333 Cholesky factorization, $A = LDL^T$, one column at a time. At iteration k , the left-
334 looking REF Cholesky factorization accesses the first $k - 1$ columns of L and D and
335 computes the k th column of L . Prior to presenting this factorization, we introduce
336 the following notation.

337 **DEFINITION 3.4.** Let $L_L^{(k)}$ and $D_L^{(k)}$ denote the k th left-looking L and D matrices
338 for $k = 0, \dots, n$. Specifically, $L_L^{(k)}$ and $D_L^{(k)}$ are the first k columns of L and D
339 augmented by I and $1/\rho^{k-1}I$, respectively.

340 Theorem 3.5 formally introduces the left-looking REF Cholesky factorization.

341 **THEOREM 3.5.** The REF Cholesky factorization can be obtained in a left-looking
342 fashion by solving Equation 6 for $k = 1, \dots, n$.

$$343 \quad (6) \quad L_L^{(k-1)} D_L^{(k-1)} \mathbf{x} = A(:, k)$$

344 Theorem 3.5 follows directly from applying Theorem 3.1 and Lemma 3.2 of [33]
345 to an SPD matrix. However, since the left-looking REF Cholesky factorization is
346 algorithmically different than traditional left-looking Cholesky factorization, below
347 we present an alternate proof of Theorem 3.5. This proof illustrates that computing
348 \mathbf{x} via the traditional approach (using a matrix-vector multiplication) does not preserve

349 integrality because it breaks the interrelation of consecutive pivots in IPGE; thus a
 350 specially-designed REF triangular solve algorithm must be used.

351 *Proof.* We prove Theorem 3.5 via a 3×3 block matrix decomposition. The REF
 352 Cholesky factorization of an SPD matrix A is given as $A = LDL^T$. Let $L_{i,j}$, $D_{i,j}$,
 353 and $A_{i,j}$, be the i,j block of L , D , and A , respectively. At iteration k , we decompose
 354 this factorization as follows:

$$355 \begin{bmatrix} L_{1,1} & 0 & 0 \\ L_{2,1} & L_{2,2} & 0 \\ L_{3,1} & L_{3,2} & L_{3,3} \end{bmatrix} \begin{bmatrix} D_{1,1} & 0 & 0 \\ 0 & D_{2,2} & 0 \\ 0 & 0 & D_{3,3} \end{bmatrix} \begin{bmatrix} L_{1,1}^T & L_{2,1}^T & L_{3,1}^T \\ 0 & L_{2,2}^T & L_{3,2}^T \\ 0 & 0 & L_{3,3}^T \end{bmatrix} = \begin{bmatrix} A_{1,1} & A_{2,1}^T & A_{3,1}^T \\ A_{2,1} & A_{2,2} & A_{3,2}^T \\ A_{3,1} & A_{3,2} & A_{3,3} \end{bmatrix}$$

356 where the first row and column of the above matrices represent rows/columns 1 to
 357 $k-1$, the second row and column represent row/column k , and the last row/column
 358 represent rows/columns $k+1$ to n , respectively.

359 Since this is a left-looking algorithm, we want to compute the k th column of L
 360 (i.e., $L_{2,2}$ and $L_{3,2}$). Via matrix multiplication, we obtain the following two equations:

$$361 \quad (7) \quad L_{2,1}D_{1,1}L_{2,1}^T + L_{2,2}D_{2,2}L_{2,2}^T = A_{2,2},$$

362

$$363 \quad (8) \quad L_{3,1}D_{1,1}L_{2,1}^T + L_{3,2}D_{2,2}L_{2,2}^T = A_{3,2}.$$

364 Since $L_{2,2} = \rho^k$ and $D_{2,2} = 1/(\rho^k \rho^{k-1})$, the expression $L_{2,2}D_{2,2}L_{2,2}^T$ is equivalent
 365 to ρ^k/ρ^{k-1} . Thus, we can solve Equation (7) for ρ^k as:

$$366 \quad \rho^k = \rho^{k-1}(A_{2,2} - L_{2,1}D_{1,1}L_{2,1}^T).$$

367 Since $A_{2,2} = a_{k,k}^0$ and $L_{2,1}D_{1,1}L_{2,1}^T = \sum_{i=1}^{k-1} \frac{l_{i,k}^2}{\rho^{i-1}\rho^i}$, the above equation becomes:

$$368 \quad (9) \quad \rho^k = \rho^{k-1}\left(a_{k,k}^0 - \sum_{i=1}^{k-1} \frac{l_{i,k}^2}{\rho^{i-1}\rho^i}\right).$$

369 Next, since $D_{2,2}L_{2,2}^T = 1/\rho^{k-1}$, Equation (8) can be solved for $L_{3,2}$ as:

$$370 \quad (10) \quad L_{3,2} = \rho^{k-1}(A_{3,2} - L_{3,1}D_{1,1}L_{2,1}^T).$$

371 Expanding the expression $L_{3,1}D_{1,1}L_{2,1}^T$, we obtain:

$$372 \quad L_{3,1}D_{1,1}L_{2,1}^T = \begin{bmatrix} \sum_{i=1}^{k-1} \frac{l_{k+1,i}l_{k,i}}{\rho^{i-1}\rho^i} \\ \vdots \\ \sum_{i=1}^{k-1} \frac{l_{n,i}l_{k,i}}{\rho^{i-1}\rho^i} \end{bmatrix}$$

373 Finally, substituting the above expression into Equation (10) and solving for an
 374 arbitrary $l_{j,k} \in L_{3,2}$, we obtain:

$$375 \quad (11) \quad l_{j,k} = \rho^{k-1}\left(a_{j,k}^0 - \sum_{i=1}^{j-1} \frac{l_{j,i}l_{k,i}}{\rho^{i-1}\rho^i}\right) \quad \text{for } j = k+1, \dots, n.$$

376 Note that Equations (9) and (11) are of the same format as Equation (8) from
 377 [33] (with $u_{i,k} = l_{i,k}^T = l_{k,i}$ due to symmetry). Thus, as argued in that paper, in order
 378 to maintain integrality, these equations must be solved with a lower triangular solve
 379 function. \square

380 Theorem 3.5 shows how to compute the left-looking REF Cholesky factorization
 381 by solving a sequence of lower trapezoidal systems (Equation 6). Similar to the
 382 discussion in Section 3.1, directly applying either a floating-point algorithm or SLIP
 383 LU’s triangular solve would be inexact or inefficient, respectively. Thus, we modify
 384 Algorithm 1 to both exploit the properties of SPD matrices and symmetry (while
 385 maintaining its integer-preserving property) as follows:

386 Symbolically, the major difference is how the nonzero pattern \mathcal{X} is computed.
 387 Computing the reach on the elimination tree (as described in Subsection 3.1) would
 388 only give the indices of the nonzeros located in the first $k - 1$ rows of the k th column
 389 of L . It is also possible to compute the reach directly as is done for LU factorization;
 390 but this is overly inefficient. Instead, as is done in other modern left-looking Cholesky
 391 factorizations, we preallocate the matrix L prior to factorization [8]. Essentially, this
 392 consists of analyzing the elimination tree to compute the entire nonzero pattern of L
 393 prior to the factorization; then, the unsorted \mathcal{X} is given as input to the left-looking
 394 factorization.

395 Numerically, we avoid half of the operations of the unsymmetric left-looking REF
 396 triangular solve. Specifically, the first $k - 1$ entries of \mathbf{x} are given as the k th row of L
 397 (which has already been computed). Thus the first $k - 1$ entries are not recomputed
 398 and are only used to update the nonzeros located in rows $k : n$ of \mathbf{x} . This modification,
 399 along with the aforementioned symbolic analysis changes, lead to Algorithm 3, an
 400 efficient method to compute the k th column of the REF Cholesky factor L .

Algorithm 3 Symmetric Sparse REF Triangular Solve (Left-Looking)

```

1: Initialization:  $x(1 : k - 1) = L(k, :)$ ,  $x(k : n) = A(k : n, :)$ 
2: sort( $\mathcal{X}$ )
3: Initialize history vector:  $h_j = 0 \forall j \in \mathcal{X}$ 
4: for  $j \in \mathcal{X}$  do
5:   if  $j < k$  then
6:     for  $i > j : i \geq k$  do
7:       if  $h_i < j - 1$  then
8:         History Update:  $x_i = \frac{x_i \rho^{j-1}}{\rho^{h_i}}$ 
9:       end if
10:      IPGE Update:  $x_i = \frac{\rho^j x_i - l_{i,j} x_j}{\rho^{j-1}}$ 
11:      History Vector Update:  $h_i = j$ 
12:    end for
13:   else
14:     if  $h_j < j - 1$  then
15:       History Update:  $x_j = \frac{x_j \rho^{j-1}}{\rho^{h_j}}$ 
16:     end if
17:   end if
18: end for

```

401 **4. Computational Complexity.** This section derives the computational com-
 402 plexity of the up-looking and left-looking REF Cholesky factorizations. Prior to
 403 deriving these complexities, we introduce the following lemma.

404 LEMMA 4.1. *Based on the best known fast Fourier transform algorithm, the cost*
 405 *of performing multiplications and divisions on two integers of bit length β is given by*
 406 $\mathcal{O}(\beta \log \beta \log \log \beta)$ [28, 41].

407 Also, given an SPD matrix $A \in \mathbb{Z}^{n \times n}$ with largest initial entry, σ , Lemma 2.2
 408 implies that the maximum bit-length required to store any entry in the REF Cholesky
 409 factorizations, denoted β_{max} , is:

$$410 \quad \beta_{max} \leq n \log \sigma.$$

411 Prior to discussing the complexities of the sparse REF Cholesky factorizations,
 412 we note that, in contrast to the dense case, there are three special considerations
 413 common to all sparse matrix factorization complexity analyses:

- 414 1. Most sparse algorithms preorder the input matrix, A . In the pre-ordering process,
 415 A is analyzed and its columns/rows are ordered to reduce fill/operations. Critically,
 416 this problem is NP-Hard; consequently, it is solved via heuristics. Most importantly,
 417 sparse factorization complexity analysis assumes that the order of the rows and
 418 columns of A is fixed (that is, the cost of symbolic preordering is never included
 419 in the complexity analysis of a sparse factorization).
 420
- 421 2. Sparse algorithms rely more heavily on manipulating intricate data structures;
 422 indeed, poor data manipulation or failure to exploit special objects like the
 423 elimination tree can cause the algorithm to have the same asymptotic complexity
 424 as a dense algorithm [19, 8, 13].
- 425 3. The number of arithmetic operations required to factorize a sparse matrix,
 426 A , does not depend *exclusively* on the matrix' dimension, n , but it heavily
 427 depends on the number and structure of nonzeros in A . For example, in
 428 floating-point arithmetic, if A is tridiagonal it requires $\mathcal{O}(n)$ arithmetic
 429 operations; however, if A is dense it requires $\mathcal{O}(n^3)$ arithmetic operations.

430 Thus, in order to better capture the amount of work, and, more precisely, in order
 431 to better distinguish between the (sparsity-specific) overhead of data structure and the
 432 necessary (arithmetic) work, the complexity of sparse matrix algorithms is typically
 433 given in terms of the number of indispensable floating-point arithmetic operations
 434 performed, denoted f [8, 13]. For the examples above $f = n$ if A is tridiagonal and
 435 $f = n^3$ if A is dense. Thus, f is a convenient term to encapsulate the number of
 436 *arithmetic* operations performed by a sparse matrix algorithm (each of which, in the
 437 floating-point case, require $\mathcal{O}(1)$ work).

438 Sparse matrix algorithms whose complexity depends solely on the cost of the arith-
 439 metic operations are called “proportional to arithmetic work” meaning that the cost
 440 of all ancillary operations is asymptotically dominated by the cost of the indispensable
 441 arithmetic work. For example, both Gilbert and Peierl’s left LU [19] and up-looking
 442 Cholesky [31] have complexities of $\mathcal{O}(f)$ meaning that the cost of the arithmetic work
 443 asymptotically dominates the cost of all other operations. Conversely, an algorithm
 444 such as right-looking factorization performs ancillary data manipulation which is not
 445 asymptotically dominated by the cost of the arithmetic operations. Lastly, though
 446 floating-point left-looking LU and up-looking Cholesky are asymptotically efficient,
 447 their rational-arithmetic counterparts are not due to their usage of ancillary greatest
 448 common divisor operations that are required to limit the bit-length growth of entries.

449 To keep with this convention, we define two new terms: I_U and I_L which denote
 450 the total number of integer-arithmetic operations required to compute the sparse up-
 451 looking and left-looking REF Cholesky factorizations, respectively. Subsection 4.1
 452 presents and proves the computational complexities of both factorizations using these
 453 terms. Then, Subsection 4.2 provides bounds on I_L and I_U to provide intuition behind
 454 the meaning of these terms.

455 **4.1. Sparse REF Cholesky Complexities.** Theorems 4.2 and 4.3 formally
 456 give the computational complexity of each factorization.

457 **THEOREM 4.2.** *The computational complexity of the sparse up-looking REF Cholesky*
 458 *factorization is*

$$459 \quad \mathcal{O}(I_U[\beta_{max} \log \beta_{max} \log \log \beta_{max}]).$$

460 **THEOREM 4.3.** *The computational complexity of the sparse left-looking REF Cholesky*
 461 *factorization is*

$$462 \quad \mathcal{O}(I_L[\beta_{max} \log \beta_{max} \log \log \beta_{max}]).$$

463 Below, we prove Theorem 4.2. Since the proof of Theorem 4.3 is very similar, it is
 464 omitted for brevity. Prior to proving Theorem 4.2, we present and prove the following
 465 Lemma.

466 **LEMMA 4.4.** *Let \hat{I}_U denote the total number of integer-arithmetic operations per-*
 467 *formed by the up-looking REF triangular solve, Algorithm 2. The computational*
 468 *complexity of this algorithm is*

$$469 \quad \mathcal{O}(\hat{I}_U[\beta_{max} \log \beta_{max} \log \log \beta_{max}]).$$

470 *Proof.* For Cholesky factorization, the nonzero pattern, \mathcal{X} , is obtained by ana-
 471 lyzing the elimination tree; thus lines 1, 2, 3, and 4 have complexities $\mathcal{O}(|A(:,k)|)$,
 472 $\mathcal{O}(|\mathcal{X}|)$, $\mathcal{O}(|\mathcal{X}| \log |\mathcal{X}|)$, and $\mathcal{O}(|\mathcal{X}|)$, respectively. Thus, the total cost of lines 1-4 is
 473 $\mathcal{O}(|\mathcal{X}| \log |\mathcal{X}|)$. Note that $|\mathcal{X}| \leq n$ because the vector is reused in between iterations
 474 and thus stays a constant size n throughout the factorization.

475 Now, we analyze lines 5-24. As argued in Subsection 3.1 and in [19, 33], each
 476 IPGE operation corresponds to an equivalent Gaussian elimination operation and the
 477 triangular solve performs only the necessary number of operations, making these lines
 478 require $\mathcal{O}(\hat{I}_U)$ total operations. Thus, the complexity for these lines is the total num-
 479 ber of operations multiplied by their cost, which is $\mathcal{O}(\hat{I}_U[\beta_{max} \log \beta_{max} \log \log \beta_{max}])$.

480 To conclude the proof, we now show that the cost $\mathcal{O}(|\mathcal{X}| \log |\mathcal{X}|)$ is dominated
 481 by the cost of operations $\mathcal{O}(\hat{I}_U[\beta_{max} \log \beta_{max} \log \log \beta_{max}])$. Since $\beta_{max} = n \log \sigma$, we
 482 have:

$$483 \quad \mathcal{O}(\hat{I}_U[\beta_{max} \log \beta_{max} \log \log \beta_{max}]) = \mathcal{O}(\hat{I}_U[n \log \sigma \log(n \log \sigma) \log \log(n \log \sigma)]).$$

484 Thus, since $|\mathcal{X}| \log |\mathcal{X}| \leq n \log n$, the complexity of Algorithm 2 is indeed:

$$485 \quad \mathcal{O}(\hat{I}_U[\beta_{max} \log \beta_{max} \log \log \beta_{max}]). \quad \square$$

486 We now prove Theorem 4.2.

487 *Proof.* This complexity follows from Lemma 4.4. Specifically, the overall factor-
 488 ization consists of repeated triangular solves; thus, the complexity of its arithmetic
 489 operations is equal to the total number of operations performed, I_U , multiplied by
 490 their cost. In addition, prior to factorization, the vectors \mathcal{X} , \mathbf{h} , and \mathbf{x} are initialized
 491 requiring $\mathcal{O}(n)$. This leads to a complexity of $\mathcal{O}(n + I_U[\beta_{max} \log \beta_{max} \log \log \beta_{max}])$.
 492 (Note that the cost of creating the elimination tree of A is asymptotically dominated
 493 by using the tree in the triangular solve [8]).

494 The proof is completed by noting that $\mathcal{O}(n) \subset \mathcal{O}(I_U)$ since the factorization must
 495 perform at least one operation on each diagonal entry (lines 6-8 in Algorithm 2). \square

496 In a bit-complexity model, these complexities mean that both the up-looking and
 497 left-looking REF Cholesky factorizations solve the SPD linear system, $A\mathbf{x} = \mathbf{b}$, in
 498 time proportional to integer-arithmetic work. To date, these are the only exact fac-
 499 torizations specifically tuned for SPD linear systems with this asymptotically efficient
 500 bound.

501 **4.2. Intuition on I_U and I_L .** The complexities discussed above depend on
 502 the terms I_U and I_L which are the total number of integer-arithmetic operations
 503 required for up-looking and left-looking factorization, respectively. However, much
 504 like the term f in floating-point arithmetic, these terms can be quite vague. In this
 505 subsection, we discuss some bounds on these terms in order to both relate them to
 506 the dimension of the matrix and gain some intuition.

507 First, note that each of these terms is lower bounded by n . Consider the case
 508 where A is diagonal. In this case, both the up-looking and left-looking algorithms must
 509 perform one multiplication and one division on each diagonal entry in rows/columns
 510 $2 : n$ (lines 6-8 of Algorithm 2 and 14-16 of Algorithm 3, respectively). Thus, for a
 511 diagonal matrix, either algorithm must perform $2(n - 1)$ operations showing that I_L
 512 and I_U are both lower bounded by $\mathcal{O}(n)$.

513 For an upper bound, consider the case where A is dense. In this case, there is
 514 no sparsity to be exploited and operations must be performed on every entry in the
 515 matrix. As a result, $\mathcal{O}(n^3)$ operations are performed.

516 This dramatic difference between the lower and upper bound provides further
 517 intuition as to why the complexity of sparse matrix algorithms are not given in terms
 518 of n . That said, for most sparse input matrices, it has been shown that the number
 519 of operations required in Cholesky factorization is in the range of $[\mathcal{O}(n^{3/2}), \mathcal{O}(n^2)]$
 520 [20, 19, 21].

521 **5. Computational Results.** This section computationally analyzes the two
 522 new factorizations and is organized as follows. Subsection 5.1 describes the speci-
 523 fications of the computational study. Subsection 5.2 compares the two sparse REF
 524 Cholesky algorithms. Subsection 5.3 compares our sparse REF Cholesky factoriza-
 525 tions to a rational-arithmetic Bunch-Parlett LDL factorization [6] as well as the un-
 526 symmetric SLIP LU factorization. Lastly, Subsection 5.4 benchmarks the relative
 527 forward error of MATLAB sparse backslash on real world SPD systems.

528 Throughout these computational tests, we use Dolan and Moré [11] performance
 529 profiles when comparing competing algorithms/approaches. Briefly, a performance
 530 profile is a tool which takes into account both the number of instances solved as well
 531 as the cost required to solve each instance. The performance of each algorithm cor-
 532 responds to a curve on a graph, where each point on the curve is what percentage of
 533 instances (y-axis) the algorithm solved within a time-multiple (x-axis) of the fastest
 534 solution time (among all algorithms) for each instance. An important property of per-
 535 formance profiles is that they are insensitive to the relative difficulty among different
 536 instances (i.e., they are not biased toward easy or hard instances). This is because
 537 given an instance, all solution times are relative to the fastest solver on that instance.
 538 The simplest way to interpret performance profiles is that the highest curve on the
 539 graph corresponds to the best performing algorithm. Lastly, as is common practice,
 540 all times in the performance profiles are shifted by 1 second, thereby precluding dra-
 541 matic yet potentially misleading results from instances with insignificant run time
 542 differences.

543 **5.1. Specifications.**

544 **5.1.1. Competitor Algorithms.** Our first set of tests evaluate the run times of
 545 the sparse REF Cholesky factorizations. To date, direct methods could exactly solve
 546 SPD linear systems in one of two ways: rational-arithmetic Bunch-Parlett LDL fac-
 547 torization or exact unsymmetric LU factorization. In [33] it was shown that the SLIP
 548 LU factorization dramatically outperforms rational LU factorization. Thus, we com-
 549 pare our sparse REF Cholesky factorizations to rational-arithmetic LDL factorization
 550 and the SLIP LU factorization. The code for SLIP LU [32] was obtained from https://github.com/clouren/SLIP_LU
 551 and is also hosted at www.suitesparse.com. However,
 552 since we could not find a readily available, commercial-quality rational-arithmetic LDL
 553 factorization, we used the GNU Multiple Precision Arithmetic (GMP) [22] Library
 554 rational-arithmetic data type as a class template argument for the Eigen [26] solver’s
 555 LDL factorization routines. Henceforth, this rational LDL factorization is referred to
 556 as Q-LDL.

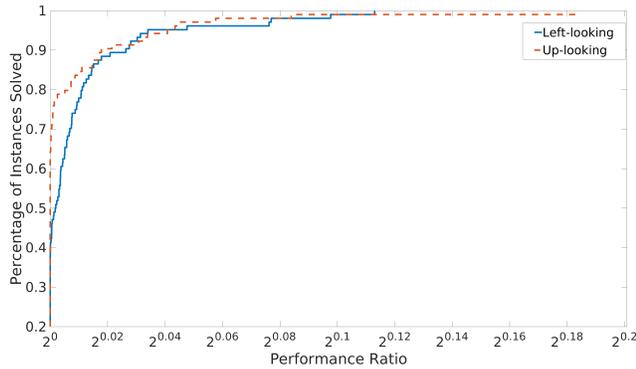
557 **5.1.2. Computing Environment.** The experiments conducted in Subsections
 558 5.2 and 5.3 measure run time and were coded in C and performed on a computing
 559 node running CentOS 7.6.1810 which has 22 GB of RAM shared by two 2.8 GHz
 560 quad core Intel Xeon 5560 processors. The experiments conducted in Subsection 5.4
 561 measure the accuracy of MATLAB sparse backslash and were performed in MATLAB
 562 R2020a on a computer running Ubuntu 18.04 with 16 GB of RAM using a 4.0 GHz
 563 Intel Core i7-8550U CPU. As mentioned in the previous subsection, all full precision
 564 integer and rational arithmetic operations were performed with the GMP library.

565 **5.1.3. Chosen Instances.** We tested our algorithms on a subset of the SuiteS-
 566 parse Matrix Collection [9], a collection of over 2800 real-world matrices arising from a
 567 vast array of applications. To obtain the subset of SPD instances from this collection,
 568 we first selected the instances which MATLAB flagged as SPD, and then verified if
 569 they were indeed SPD using our sparse REF Cholesky factorizations. We kept a total
 570 of 104 instances: all those which were both verified to be SPD and could be factor-
 571 ized within 24 hours. Appendix A includes the comprehensive results along with the
 572 indices of each selected matrix from this collection.

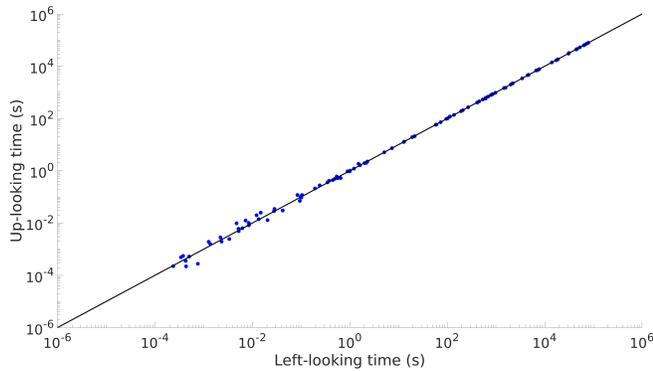
573 **5.1.4. Symbolic Analysis.** Cholesky factorization algorithms are almost al-
 574 ways preceded by a symbolic analysis phase in which a permutation matrix P
 575 is chosen so that the Cholesky factorization of the matrix PAP^T requires less work to
 576 compute than the direct factorization of A . Finding a permutation which minimizes
 577 the number of nonzeros in the Cholesky factor is NP-Complete [46]; similarly, finding
 578 a permutation which minimizes the total work required to perform a (floating-point)
 579 Cholesky factorization is NP-Hard [34]. Exact factorization differs from these prob-
 580 lems in that the total work to be performed is not only a function of the number of
 581 operations but also the cost of each operation (which is not constant). As a result,
 582 it is nearly certain that finding a permutation which minimizes the total work for
 583 these exact factorizations is also NP-Hard. Thus, for the symbolic analysis phase, we
 584 utilize the approximate minimum degree (AMD) ordering [1, 2], the state-of-the-art
 585 approach to efficiently reduce fill-in in the Cholesky factor. Specifically, we preordered
 586 all matrices with AMD prior to executing the exact REF or LDL factorizations.

587 **5.2. Comparison of Sparse REF Cholesky Algorithms.** This subsection
 588 compares the performance of the left-looking and up-looking REF Cholesky factoriza-
 589 tions. Across all 104 instances the algorithms are essentially identical; the up-looking
 590 factorization’s average run-time was about 0.5% smaller than the left-looking’s run
 591 time, while the left-looking factorization’s geometric-mean run-time was about 1%

592 smaller than that of up-looking. Figure 3 confirms graphically that both algorithms
 593 perform almost identically. Specifically, Figure 3a is a performance profile showing
 594 that, for all of the instances, both algorithms' run-times are within 7% of each other.
 595 Indeed, even though the profile appears to show that up-looking factorization is supe-
 596 rior, notice that the magnitude of differences between the two algorithms (i.e., the x
 597 axis) is practically insignificant. Figure 3b further illustrates this via a scatter plot of
 598 the run-time of both algorithms for each instance. Therein, any dot lying below the
 599 line corresponds to a matrix in which left-looking required more run-time, and vice
 600 versa. Note how the only minor performance differences occurred in instances where
 601 both algorithms ran in less than 1 second.



(a) Left and Up-looking Perform Similarly



(b) Our sparse REF Cholesky factorizations are Nearly Identical

Fig. 3: Comparison of Left-looking and Up-looking Cholesky

602 **5.3. Sparse REF Cholesky vs Alternate Direct Methods.** This subsection
 603 compares the performance of the sparse REF Cholesky factorizations to both
 604 the unsymmetric SLIP LU factorization and the rational LDL factorization (Q-LDL).
 605 Since the up-looking and left-looking REF Cholesky factorizations have similar per-
 606 formance; for simplicity, this subsection only compares the left-looking algorithm to
 607 the alternate direct methods.

608 First, we compare left-looking REF Cholesky factorization to the SLIP LU fac-

609 torization. Since SLIP LU is an unsymmetric algorithm, we would expect it to require
 610 about double the run time of the left-looking REF Cholesky factorization. Across the
 611 104 instances, REF Cholesky was faster than SLIP LU for 85% of the instances with
 612 an average and geometric mean run time 1.70 and 1.86 times smaller, respectively.
 613 Note that these results are slightly biased in favor of SLIP LU as it could not factor
 614 6 of the instances within a 24 hour time period; thus its run time was set to this upper
 615 bound. Graphically, we see these results via figure 4a which gives a performance
 616 profile comparing these two algorithms which shows that left-looking Cholesky fac-
 617 torization clearly outperforms SLIP LU. Fix. Similarly, Figure 4b, which is a scatter
 618 plot of the algorithm’s run-times, illustrates that Cholesky left-looking factorization
 619 outperforms SLIP LU on over 80% of the instances.

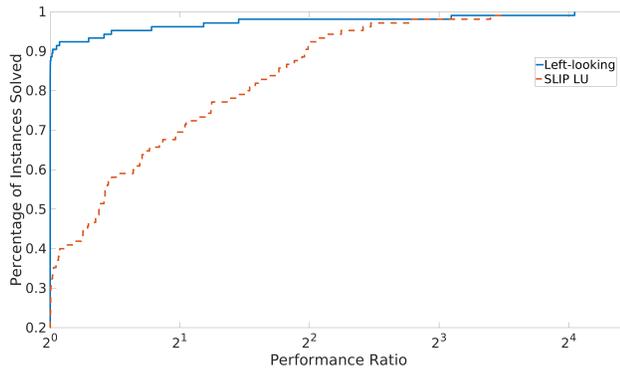
620 While REF Cholesky outperforms SLIP LU, its dominance is *not uniform* as ex-
 621 pected. Specifically, SLIP LU outperforms REF Cholesky in about 15% of the test in-
 622 stances, a counter-intuitive result which would not occur in floating-point arithmetic.
 623 This behavior can be attributed to the interplay between the fundamental nature of
 624 SPD instances and the properties of integer-arithmetic factorization algorithms:

- 625 1. The majority of SPD instances are diagonally dominant [21] meaning that
 626 the magnitudes of their diagonal entries are larger than the sum of all other
 627 entries in their respective row/column.
- 628 2. Lourenco et al. [32] showed that IPGE-based algorithms, such as SLIP LU
 629 and REF Cholesky, perform best when small pivots are selected (specifically,
 630 in the unsymmetric case, selecting small pivots is about 4 times faster than
 631 selecting large pivots).
- 632 3. Cholesky is bound to select the (large in magnitude) diagonal entries as pivots,
 633 while SLIP LU has the freedom to select (relatively smaller) off-diagonal
 634 pivots.

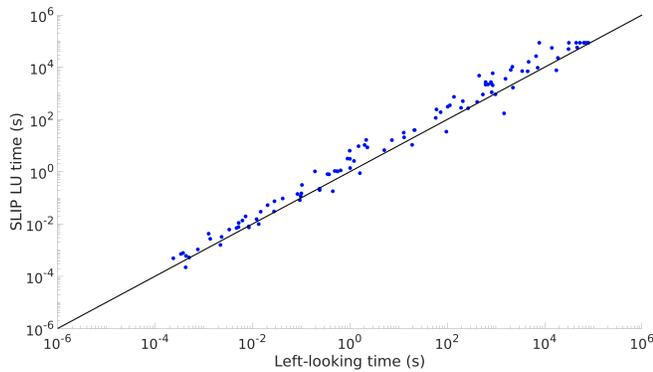
635 Consequently, even though SLIP LU generally performs about twice as many op-
 636 erations as REF Cholesky, in a small subset of instances, the smaller cost of the
 637 operations in SLIP LU allows it to negate this disadvantage. Indeed, this result pro-
 638 vides motivation for the eventual development of an integral LDL factorization in the
 639 vein of [12] which would allow symmetric, off diagonal (thus smaller in magnitude)
 640 pivoting.

641 We next compare left-looking REF Cholesky factorization to rational-arithmetic
 642 LDL factorization. Across the 104 instances, REF Cholesky is faster than rational
 643 LDL on 83% of the tested instances, and its average and geometric mean run-times are
 644 2.04 and 2.22 times smaller, respectively. Moreover, like SLIP LU, these results are
 645 slightly biased in favor of rational LDL, as this method could not factorize 8 instances
 646 within 24 hours; thus the run times for these matrices were set to this upper bound.
 647 These result support the following important observations:

- 648 1. Rational-arithmetic is incredibly slow. Notably, this can be almost exclu-
 649 sively attributed to the rational-arithmetic operations (and, specifically to
 650 their intrinsic GCD operations) because both REF Cholesky and rational
 651 LDL return matrices with the same nonzero pattern (both utilize AMD as a
 652 preordering and pivot along the diagonal). Conversely, the savings of REF
 653 Cholesky over rational LDL can be almost entirely attributed to its use of
 654 fast integer-arithmetic.
- 655 2. Sparse REF Cholesky factorization outperforms Q-LDL for over 80% of in-
 656 stances including *every* instance that requires more than two minutes except
 657 for one matrix (index 43). In contrast, the vast majority of the instances
 658 where Q-LDL is faster than left-looking REF Cholesky have insignificant run



(a) Left-looking Cholesky Outperforms SLIP LU



(b) Left-looking Cholesky Outperforms SLIP LU

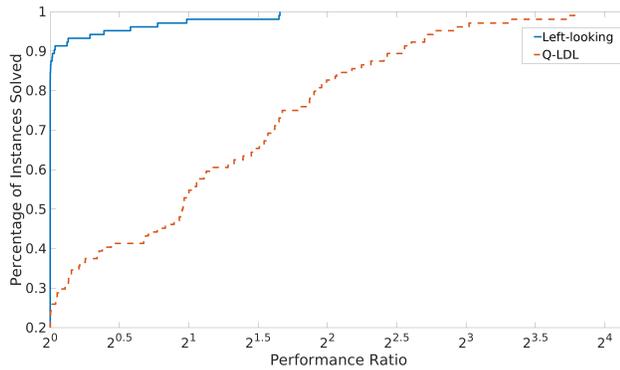
Fig. 4: Comparison of Left-looking Cholesky with SLIP LU

659 times for both methods (less than 1 second).

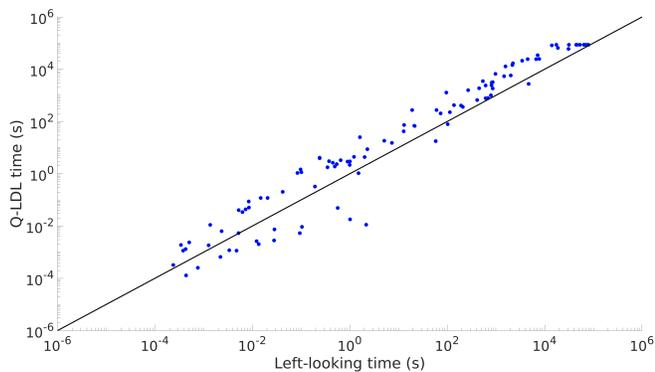
660 Graphically, the performance profile in Figure 5a and the scatter plot of the run-
 661 times in Figure 5b depict the superiority of REF Cholesky over rational LDL.

662 We conclude this subsection with the following interesting note: Rational LDL
 663 factorization performs so poorly in general that it is slower than even the unsymmetric
 664 SLIP LU factorization. Altogether, these results present compelling evidence of the
 665 superiority of our sparse REF Cholesky factorizations over both the best available
 666 exact LU factorization (SLIP LU) and rational-arithmetic LDL factorization.

667 **5.4. Relative Forward Error of MATLAB Sparse Backslash.** As dis-
 668 cussed in the introduction, Cholesky factorization is known to be normwise backward
 669 stable [45] meaning that the relative residual $\|Ax - \mathbf{b}\| / (\|A\| \|\mathbf{x}\|)$ is guaranteed to
 670 be in the neighborhood of machine precision. That said, a small relative residual does
 671 not necessarily mean that the solution \mathbf{x} is close to the exact solution of the system
 672 [21]. In this section, we use our sparse REF Cholesky factorizations to benchmark
 673 the relative forward error of MATLAB backslash on our set of SPD matrices. For
 674 this purpose, we solve the linear system $A\mathbf{x} = \mathbf{1}$ in both MATLAB and with our
 675 sparse REF Cholesky factorization. Our exact solution is then converted to double



(a) Left-looking Cholesky Outperforms Q-LDL



(b) Left-looking Cholesky Outperforms Q-LDL in Run Time

Fig. 5: Comparison of Left-looking Cholesky with Q-LDL

676 precision (note that internally, this conversion is done in higher precision; thus this
 677 solution from REF Cholesky is accurate to machine precision). Finally, we compare
 678 this rounded-to-double exact solution, $\mathbf{x}^{(e)}$, to the solution returned by MATLAB, \mathbf{x} ,
 679 by computing the relative forward error $\|\mathbf{x}^{(e)} - \mathbf{x}\|/\|\mathbf{x}^{(e)}\|$ using the 2-norm within
 680 MATLAB.

681 Table 1 gives the relative forward error of MATLAB sparse backslash on the
 682 tested linear systems. Overall, the Cholesky routines produce nearly exact solutions
 683 (relative forward error less than 10^{-12}) for 69% of tested instances and solutions with
 684 6 to 12 digits, 2-5, or fewer than 2 digits of precision for 23%, 5% and 3% of instances,
 685 respectively. While these results may seem troubling, it is important to contextualize
 686 them. Specifically, it is important to note that: (1) the majority of the instances
 687 in which the MATLAB solution was not accurate have a smallest singular value in
 688 the neighborhood of machine precision (as approximately computed with MATLAB
 689 dense SVD)—thereby meeting Demmel’s [10] criterion for floating-point Cholesky to
 690 fail; and (2) unlike the LU routines, the Cholesky routines within MATLAB sparse
 691 backslash do not perform any iterative refinement to fine tune the solutions.

692 From these results we draw two conclusions. First, enabling iterative refinement

693 within MATLAB Cholesky routines would likely increase the accuracy of the solutions
 694 returned. Second, exact methods are necessary for a modest subset of real world
 695 problems in which finite-precision Cholesky factorizations are not adequate.

Table 1: Relative Forward Error of MATLAB on SPD Linear Systems

Relative Forward Error Threshold	Percentage of Instances
$\leq 10^{-12}$	69.23%
$\leq 10^{-6}$	92.31%
$\leq 10^{-2}$	97.12%

696 **6. Conclusion.** This paper derives two new factorization algorithms which ex-
 697 actly solve sparse SPD linear systems exclusively using integer-arithmetic. Applied to
 698 a sparse matrix with a fixed row and column ordering, both derived algorithms have
 699 the property that their worst-case time complexity is proportional to the cost of the
 700 necessary integer-arithmetic work. Our derived sparse REF Cholesky factorizations
 701 require no square roots, thus they can be applied to any SPD linear system.

702 Computationally, we show that the derived algorithms dramatically outperform
 703 both exact LU factorization and rational LDL factorization. Specifically, REF Cholesky
 704 outperforms the unsymmetric SLIP LU factorization and the rational LDL factoriza-
 705 tion on average by a factor of 1.7 and 2.04, respectively. Surprisingly, this shows
 706 that rational LDL is even outperformed by SLIP LU. Interestingly, SLIP LU outper-
 707 formed REF Cholesky on 15% of instances despite requiring twice the operations. We
 708 attribute this counter-intuitive result to the complex interplay between the number
 709 of operations performed and the cost of these operations. Specifically, unlike REF
 710 Cholesky which must select diagonal pivot elements (typically large in SPD matrice-
 711 ces), SLIP LU is able to select pivots of small bit-length thereby reducing the cost of
 712 the exact operations. Additionally, we measured the relative forward error of MAT-
 713 LAB sparse backslash, showing that, as expected from the literature, it produced
 714 solutions of low accuracy on those instances which are highly ill-conditioned and/or
 715 have extremely small singular values.

716 Altogether, these results (1) show REF Cholesky is the best among competing
 717 factorization approaches for exactly solving SPD linear systems, (2) provide evidence
 718 that fixed precision floating-point Cholesky factorization can fail on real world highly
 719 ill-conditioned systems which do not meet the stability bounds proposed by [45, 10,
 720 24], and (3) illustrate that commercial Cholesky routines would likely benefit from
 721 iterative refinement—a technique that is currently used in modern LU factorization
 722 routines.

723 In conclusion, the derived sparse REF Cholesky factorizations provide a robust
 724 framework to exactly solve SPD linear systems. The code associated with these
 725 factorizations is hosted at https://github.com/clouren/IP_Chol.

726 **Acknowledgments.** The authors would like to thank the reviewers for their
 727 very helpful comments which significantly improved the quality of this paper. The
 728 work is partially supported by the National Science Foundation under Grant No.
 729 OAC-1835499. In addition, the first author was also partially supported by the Texas
 730 A&M University’s Graduate Merit Fellowship and the United States Naval Academy
 731 Junior NARC.

732

REFERENCES

- 733 [1] P. R. AMESTOY, T. A. DAVIS, AND I. S. DUFF, *An approximate minimum degree ordering*
734 *algorithm*, SIAM Journal on Matrix Analysis and Applications, 17 (1996), pp. 886–905.
- 735 [2] P. R. AMESTOY, T. A. DAVIS, AND I. S. DUFF, *Algorithm 837: Amd, an approximate mini-*
736 *imum degree ordering algorithm*, ACM Transactions on Mathematical Software (TOMS),
737 30 (2004), pp. 381–388.
- 738 [3] R. E. BANK AND R. K. SMITH, *General sparse elimination requires no permanent integer*
739 *storage*, SIAM journal on scientific and statistical computing, 8 (1987), pp. 574–584.
- 740 [4] E. H. BAREISS, *Sylvester’s identity and multistep integer-preserving gaussian elimination*,
741 *Mathematics of computation*, 22 (1968), pp. 565–578.
- 742 [5] E. H. BAREISS, *Computational solutions of matrix problems over an integral domain*, IMA
743 *Journal of Applied Mathematics*, 10 (1972), pp. 68–104.
- 744 [6] J. R. BUNCH AND B. N. PARLETT, *Direct methods for solving symmetric indefinite systems of*
745 *linear equations*, SIAM Journal on Numerical Analysis, 8 (1971), pp. 639–655.
- 746 [7] Y. CHEN, T. A. DAVIS, W. W. HAGER, AND S. RAJAMANICKAM, *Algorithm 887: Cholmod, su-*
747 *pernodal sparse cholesky factorization and update/downdate*, ACM Transactions on Math-
748 *ematical Software (TOMS)*, 35 (2008), p. 22.
- 749 [8] T. A. DAVIS, *Direct methods for sparse linear systems*, SIAM, 2006.
- 750 [9] T. A. DAVIS AND Y. HU, *The university of florida sparse matrix collection*, ACM Transactions
751 *on Mathematical Software (TOMS)*, 38 (2011), p. 1.
- 752 [10] J. DEMMEL, *On floating point errors in Cholesky*, University of Tennessee. Computer Science
753 *Department*, 1989.
- 754 [11] E. D. DOLAN AND J. J. MORÉ, *Benchmarking optimization software with performance profiles*,
755 *Mathematical programming*, 91 (2002), pp. 201–213.
- 756 [12] I. S. DUFF, *Ma57—a code for the solution of sparse symmetric definite and indefinite systems*,
757 *ACM Transactions on Mathematical Software (TOMS)*, 30 (2004), pp. 118–144.
- 758 [13] I. S. DUFF, A. M. ERISMAN, AND J. K. REID, *Direct methods for sparse matrices*, Oxford
759 *University Press*, 2017.
- 760 [14] J. EDMONDS, *Systems of distinct representatives and linear algebra*, *J. Res. Nat. Bur. Standards*
761 *Sect. B*, 71 (1967), pp. 241–245.
- 762 [15] S. C. EISENSTAT, M. H. SCHULTZ, AND A. H. SHERMAN, *Algorithms and data structures for*
763 *sparse symmetric gaussian elimination*, SIAM Journal on Scientific and Statistical Com-
764 *puting*, 2 (1981), pp. 225–237.
- 765 [16] A. R. ESCOBEDO AND E. MORENO-CENTENO, *Roundoff-error-free algorithms for solving linear*
766 *systems via cholesky and lu factorizations*, *INFORMS Journal on Computing*, 27 (2015),
767 pp. 677–689.
- 768 [17] A. R. ESCOBEDO, E. MORENO-CENTENO, AND C. LOURENCO, *Solution of dense linear systems*
769 *via roundoff-error-free factorization algorithms: Theoretical connections and computa-*
770 *tional comparisons*, *ACM Transactions on Mathematical Software*, (2018).
- 771 [18] A. GEORGE AND J. W. LIU, *Computer solution of large sparse positive definite*, Prentice Hall
772 *Professional Technical Reference*, 1981.
- 773 [19] J. R. GILBERT AND T. PEIERLS, *Sparse partial pivoting in time proportional to arithmetic*
774 *operations*, SIAM Journal on Scientific and Statistical Computing, 9 (1988), pp. 862–874.
- 775 [20] J. R. GILBERT AND R. SCHREIBER, *Nested dissection with partial pivoting*, in *Sparse Matrix*
776 *Symposium 1982: Program and Abstracts*, 1982, p. 61.
- 777 [21] G. H. GOLUB AND C. F. VAN LOAN, *Matrix computations*, vol. 3, JHU Press, 2012.
- 778 [22] T. GRANLUND ET AL., *GNU MP 6.0 Multiple Precision Arithmetic Library*, Samurai Media
779 *Limited*, 2015.
- 780 [23] J. HADAMARD, *Résolution d’une question relative aux déterminants*, *Bull. sci. math*, 17 (1893),
781 pp. 240–246.
- 782 [24] N. J. HIGHAM, *Accuracy and stability of numerical algorithms*, vol. 80, Siam, 2002.
- 783 [25] R. A. HORN AND C. R. JOHNSON, *Matrix analysis*, Cambridge university press, 2012.
- 784 [26] B. JACOB, G. GUENNEBAUD, ET AL., *Eigen: C++ template library for linear algebra*, 2013.
- 785 [27] A. KIELBASIŃSKI, *A note on rounding-error analysis of cholesky factorization*, *Linear Algebra*
786 *and its Applications*, 88 (1987), pp. 487–494.
- 787 [28] D. E. KNUTH, *The art of computer programming: sorting and searching*, vol. 3, Pearson Edu-
788 *cation*, 1998.
- 789 [29] H. R. LEE AND B. D. SAUNDERS, *Fraction free gaussian elimination for sparse matrices*, *Journal*
790 *of symbolic computation*, 19 (1995), pp. 393–402.
- 791 [30] J. W. LIU, *A compact row storage scheme for cholesky factors using elimination trees*, *ACM*
792 *Transactions on Mathematical Software (TOMS)*, 12 (1986), pp. 127–148.

- 793 [31] J. W. LIU, *A generalized envelope method for sparse factorization by rows*, ACM Transactions
794 on Mathematical Software (TOMS), 17 (1991), pp. 112–129.
- 795 [32] C. LOURENCO, J. CHEN, E. MORENO-CENTENO, AND T. DAVIS, *Algorithm 1xxx: Slip lu, exactly*
796 *solving sparse linear systems via a sparse left-looking integer-preserving lu factorization*.
797 Submitted ACM Transactions on Mathematical Software, 2020.
- 798 [33] C. LOURENCO, A. R. ESCOBEDO, E. MORENO-CENTENO, AND T. A. DAVIS, *Exact solution of*
799 *sparse linear systems via left-looking roundoff-error-free lu factorization in time propor-*
800 *tional to arithmetic work*, SIAM Journal on Matrix Analysis and Applications, 40 (2019),
801 pp. 609–638.
- 802 [34] R. LUCE AND E. G. NG, *On the minimum flops problem in the sparse cholesky factorization*,
803 SIAM Journal on Matrix Analysis and Applications, 35 (2014), pp. 1–21.
- 804 [35] J. MEINGUET, *Refined error analyses of cholesky factorization*, SIAM journal on numerical
805 analysis, 20 (1983), pp. 1243–1250.
- 806 [36] R. M. MONTANTE-PARDO AND M. A. MÉNDEZ-CAVAZOS, *Un método numérico para cálculo*
807 *matricial*, Revista Técnico-Científica de Divulgación, 2 (1977), pp. 1–24.
- 808 [37] T. MUIR AND W. H. METZLER, *A Treatise on the Theory of Determinants*, Courier Corporation,
809 2003.
- 810 [38] Y. U. D. OF COMPUTER SCIENCE, S. EISENSTAT, M. SCHULTZ, AND A. SHERMAN, *Efficient*
811 *implementation of sparse symmetric Gaussian elimination*, 1975.
- 812 [39] S. PARTER, *The use of linear graphs in gauss elimination*, SIAM review, 3 (1961), pp. 119–130.
- 813 [40] D. ROSE, G. WHITTEN, A. SHERMAN, AND R. E. TARJAN, *Algorithms and software for in-core*
814 *factorization of sparse symmetric positive definite matrices*, Computers & Structures, 11
815 (1980), pp. 597–608.
- 816 [41] A. SCHÖNHAGE AND V. STRASSEN, *Schnelle multiplikation grosser zahlen*, Computing, 7 (1971),
817 pp. 281–292.
- 818 [42] A. SCHRIJVER, *Theory of integer and linear programming*, 1986.
- 819 [43] J. STEIN, *Computational problems associated with racah algebra*, Journal of Computational
820 Physics, 1 (1967), pp. 397–405.
- 821 [44] J.-G. SUN, *Componentwise perturbation bounds for some matrix decompositions*, BIT Numerical
822 Mathematics, 32 (1992), pp. 702–714.
- 823 [45] J. H. WILKINSON, *A priori error analysis of algebraic processes*, in Intern. Congress Math,
824 vol. 19, 1968, pp. 629–639.
- 825 [46] M. YANNAKAKIS, *Computing the minimum fill-in is np-complete*, SIAM Journal on Algebraic
826 Discrete Methods, 2 (1981), pp. 77–79.

827 Appendix A. Comprehensive Computational Results.

828 Tables 2, 3, and 4 present exhaustive computational results for all 104 tested
829 instances sorted based on their run times. The first three columns of each table give
830 the name (i.e., matrix index from SuiteSparse collection), dimension, and number of
831 nonzeros of each matrix. Column 4 gives the run time of the left-looking Cholesky
832 factorization. Finally columns 5, 6, and 7 give the run times of up-looking Cholesky,
833 SLIP LU, and Q-LDL relative to left-looking Cholesky, respectively. Note that **N/A**
834 in relative run times indicates that the algorithm could not factorize this instance
835 within 24 hours. Lastly, any matrix whose solution had less than 2 digits of precision
836 is bolded and colored **red** and any instance whose solution had between 2 and 6 digits
837 of precision is italicized and colored *blue*.

Matrix Index	n	nnz	Left-Looking (hr)	Relative Run Time		
				Up-Looking	SLIP LU	Q-LDL
888	9801	87025	21.89	1	N/A	N/A
1330	7102	340200	20.14	0.99	N/A	N/A
760	4704	104756	18.50	1	N/A	N/A
887	9604	85264	17.68	0.99	N/A	N/A
50	4410	219024	14.69	1	N/A	N/A
46	3562	159910	12.92	1	1.23	N/A
791	8205	125567	12.28	0.99	N/A	N/A
35	2003	83883	8.78	1.01	N/A	N/A

Matrix Index	n	nnz	Left-Looking (hr)	Relative Run Time		
				Up-Looking	SLIP LU	Q-LDL
360	4515	97707	8.57	0.99	1.61	1.93
413	6867	98671	5.20	0.99	1.22	3.51
2211	2000	41906	4.83	0.98	0.44	N/A
759	2910	174296	3.92	0.99	3.93	5.85
1214	4098	72356	2.16	1	N/A	3.17
758	2146	72250	2.00	0.99	1.33	4.74
440	3363	99471	1.85	1.02	3.97	3.66
43	3600	26600	1.31	0.99	3.38	0.58
36	1806	63454	1.25	1	1.56	5.39

Table 2: Results for Instances Greater than 1 Hour

Matrix Index	n	nnz	Left-Looking (min)	Relative Run Time		
				Up-Looking	SLIP LU	Q-LDL
359	1440	44998	57.61	1	2.05	6.09
357	726	34518	37.09	1	0.75	7.58
757	1824	39208	36.34	1	4.74	6.60
<i>411</i>	2568	75628	33.60	1	3.89	2.85
1437	19998	99982	26.24	0.95	2.31	8.12
889	9801	87025	24.47	1	0.12	3.65
31	1083	18437	16.26	0.99	0.95	6.79
228	1919	32399	14.33	1	2.36	2.16
48	1922	30336	14.28	1	6.87	3.74
34	1473	34241	13.70	1	1.34	3.67
33	1473	34241	13.66	1	1.36	3.77
427	1821	52685	13.58	0.99	2.87	2.97
1331	7102	170134	13.19	1	3.40	1.28
408	2548	57308	11.39	1	3.21	1.16
1439	19994	79966	10.27	0.93	4.41	3.89
407	2410	54840	10.20	1.01	3.54	1.26
49	1224	56126	9.00	1	1.68	6.51
30	1074	12960	7.51	1	10.54	4.19
1911	1282	30644	6.84	1	1.14	1.61
339	588	21418	4.48	1	1	5.90
430	1733	22189	3.49	1	2.38	1.71
219	960	15844	3.22	0.99	1.44	2.10
<i>358</i>	1050	26198	2.30	1	5.36	3.08
223	729	4617	1.89	1.06	3.06	2.01
67	1473	19659	1.71	0.99	3.01	0.76
2210	700	12654	1.61	1	0.36	13.32
32	1086	22070	1.22	0.98	2.60	2.80
1328	500	28726	1.01	0.98	4.03	4.58

Table 3: Results for Instances Between 1 Min and 1 Hour

Matrix Index	n	nnz	Left-Looking (s)	Relative Run Time		
				Up-Looking	SLIP LU	Q-LDL
221	468	5172	58.150	1	1.99	0.31
28	420	7860	21.426	0.98	1.83	3.17
29	420	7860	21.025	0.99	1.86	3.23
2209	500	8478	19.073	1.01	0.56	14.42
62	420	7252	12.991	1	1.61	5.64
41	817	6853	12.838	0.98	2.47	3.31
222	675	3255	7.336	1	2.21	2.07
229	362	5786	5.094	1	1.33	3.62
929	415	2779	2.289	0.99	3.74	3.81
1401	2541	7361	2.173	0.92	7.65	0.01
4	685	3249	2.010	0.98	5.26	2.17
2208	300	4678	1.616	1	0.55	15.65
1	1138	4054	1.516	1.22	6.31	0.70
42	485	3135	1.219	1	2.12	3.61
75	15439	15439	1.014	1	1.36	0.02
875	306	2018	1.004	1	3.11	2.98
3	662	2474	0.991	0.96	6.40	2.20
876	306	2018	0.898	1.05	3.54	3.28
26	132	3648	0.649	0.81	1.73	5.12
1425	10001	49999	0.567	0.91	1.82	0.09
878	289	1377	0.538	1.12	1.97	4.38
206	147	2449	0.492	0.99	2.18	3.83
159	900	7744	0.450	0.98	0.40	5.86
27	153	2423	0.375	1.11	2.14	8.02
207	147	2441	0.348	1.03	2.32	5.13
2206	199	2873	0.241	1.12	0.93	17.23
2207	200	2890	0.240	1.16	0.83	16.42
2	494	1666	0.193	1.09	5.34	1.69
315	416	2312	0.104	1.13	3.02	0.09
1506	124	12068	0.101	0.94	1.46	11.35
2205	150	2040	0.097	1.02	1.20	14.99
73	3134	3134	0.094	0.76	0.88	0.06
24	66	4356	0.084	1.41	1.68	12.71
44	138	696	0.042	0.71	2.25	4.82
218	957	4137	0.028	1.22	2.64	0.26
76	1922	1922	0.028	1.04	1.08	0.10
220	100	594	0.021	0.63	2.56	5.74
23	48	400	0.015	1.70	2.02	7.97
63	1074	1074	0.014	1.04	0.74	0.15
66	1473	1473	0.012	1.64	1.26	0.22
873	48	306	0.009	0.96	0.87	5.97
877	289	1377	0.008	1.21	0.96	10.29
25	112	640	0.007	1.70	2.69	5.92
874	48	306	0.006	1.01	2.18	5.44
872	48	306	0.005	0.93	2.09	7.72
71	3600	3600	0.005	1.15	1.49	1.03
69	817	817	0.005	2.07	1.48	0.24
70	485	485	0.003	0.73	1.83	0.35

Matrix Index	n	nnz	Left-Looking (s)	Relative Run Time		
				Up-Looking	SLIP LU	Q-LDL
1939	3240	3240	0.002	0.83	1.39	2.74
61	420	420	0.002	1.28	0.72	0.30
436	27	279	0.001	1.17	1.98	8.15
217	237	1017	0.001	1.52	3.37	1.44
60	153	153	0.001	0.36	1.41	0.34
1438	18	82	0.001	1.04	1.03	4.66
57	66	66	0.000	0.50	1.37	0.30
1440	14	46	0.000	0.84	0.52	3.06
2204	20	158	0.000	1.45	2.05	3.01
2203	19	147	0.000	1.42	2.11	5.54
72	138	138	0.000	0.94	2.06	1.36

Table 4: Results for Instances Less than 1 Minute